Machine learning for discriminating quantum measurement trajectories and improving readout

Easwar Magesan, Jay M. Gambetta, A.D. Córcoles, and Jerry M. Chow
IBM T.J. Watson Research Center, Yorktown Heights, NY 10598, USA

High-fidelity measurements are important for the implementation of quantum information protocols. Current methods for classifying measurement trajectories in superconducting qubit systems produce fidelities systematically lower than those predicted by experimental parameters. Here, we place current classification methods within the framework of machine learning (ML) algorithms and improve on them by investigating more sophisticated ML approaches. We find that non-linear algorithms and clustering methods produce significantly higher assignment fidelities that help close the gap to the fidelity achievable under ideal noise conditions. Clustering methods group trajectories into natural subsets within the data, which allows for the diagnosis of systematic errors. We find large clusters in the data associated with T_1 processes and show these are the main source of discrepancy between our experimental and achievable fidelities. These error diagnosis techniques help provide a path forward to improve qubit measurements.

arXiv:1411.4994v2 [quant-ph] 24 Nov 2014

The ability to perform accurate measurements is important for maximizing the information one can extract from a physical system. This is especially true in experimental quantum information processing since quantum systems are highly susceptible to noise effects and error rates of quantum operations and measurements must be small for fault-tolerant quantum computation to be a reality [1]. Our goal is to provide methods for diagnosing measurement errors and increasing fidelities by using classification and clustering algorithms in machine learning (ML). While we apply our methods in a superconducting qubit measurement system, we anticipate that the generality of these techniques can be useful in a broader class of systems.

Superconducting quantum bits (qubits) have become a promising candidate for building a fault-tolerant quantum computer due to their long coherence times [2–4], high-fidelity multi-qubit gate operations [5], and the abilitiy to perform single-shot measurements [6–9], in a circuit QED architecture [10]. Remarkable progress has been made in reducing error-rates of these operations however considerable work is still needed to implement fault-tolerant quantum computation in large networks of qubits. In circuit quantum electrodynamics (cQED) a superconducting anharmonic oscillator, such as a transmon [11], is coupled to a resonator, producing a state-dependent shift of the resonator frequency. This allows for qubit measurements by driving the resonator and recording the output trajectory [1] in phase (I-Q) space. In practice there are significant sources of random noise and systematic effects, such as T_1 processes where the system spontaneously decays to its ground state, that can make single-shot trajectories appear complex and difficult to distinguish.

Our experimental system is a single qubit (Q4) in a planar lattice of four superconducting qubits [13]. We show current methods for assigning outcomes to measurement trajectories in this system produce assignment fidelities (defined below in Eq. (29)) that are much lower than the predicted achievable value derived under ideal noise conditions . We utilize various ML algorithms to obtain deeper insight into the behavior of the trajectories and bring fidelities up closer to our expected values. We find a total increase in assignment fidelity from 0.9586 using current methods to 0.9821 ($\sim 2.4\%$ increase) using non-linear ML classifiers. The strong performance of non-linear classifiers indicates systematic effects, such as heating and T_1 events, could be a significant source of error in our measurements.

To verify this, we use ML clustering methods to group the data into naturally occurring subsets. We find a large cluster consisting of T_1 events whose size is consistent with the experimentally measured T_1 time. Replacing this cluster with random non-T_1 events gives assignment fidelities approaching 0.995 which is much closer to the achievable value of 0.9999. Going to higher orders we find a much small cluster corresponding to heating of the ground state into the excited state. Before moving on, let us make a few points about using ML methods for trajectory discrimination and improving measurements. First, the methods we present here can be useful in a much broader context. Any measurement scheme that produces patterns in a geometric space can potentially benefit from more advanced ML methods. Investigating the applicability to different systems will depend on the details of each situation. Second, these methods are applicable even if we are trying to improve higher fidelity measurements than that of this paper. The key is that these methods can be tailored according to the types of noise present. Third, ML methods have also been applied to other problems in quantum information such as phase estimation [14] and asymptotic state estimation [15].

The standard metric for characterizing how well a single-qubit measurement assigns outcomes is the assignment fidelity

$$\mathcal{F}_a = 1 - \left(\mathbb{P}\left[0|1\right] + \mathbb{P}\left[1|0\right]\right)/2. \tag{1}$$

Here $\mathbb{P}[0|1]$ ($\mathbb{P}[1|0]$) is the probability of obtaining out-

come "0" ("1") given the system was prepared in $|1\rangle$ ($|0\rangle$). Hence $\mathcal{F}_a \in [0,1]$ and ideally $\mathcal{F}_a \sim 1$. Our measurement framework is the dispersive limit of cQED, where observing the resonator output voltage provides a quantum non-demolition qubit measurement. For outcome "0" (1) the voltage leaving the cavity represents a coherent state I-Q trajectory, and single-shot trajectories are obtained by amplifying the cavity signal. The main parameters of our system are given in [16] and complete experimental information is given in Ref. [13].

Our data consists of 51200 single-shot trajectories (shots), half initially prepared in $|0\rangle$ and the other half in $|1\rangle$ (denote these classes by C_0 and C_1). The trajectories are time-ordered and the first half is used as a training set to predict the second half for classification. The mean trajectories for each class, denoted $\mu_0(t)$ and $\mu_1(t)$, as well as examples of single-shot trajectories, are given in Fig. 1. We see that each shot can be complicated but there are enough shots to ensure smooth, well-separated means. The total measurement time T is 2.6μs and $[0,T]$ is discretized into 163 time-points so trajectories are represented by vectors $x \in \mathbb{R}^{326}$ (let $M = 326$) where the first (last) 163 entries correspond to the real (imaginary) parts of the trajectory.

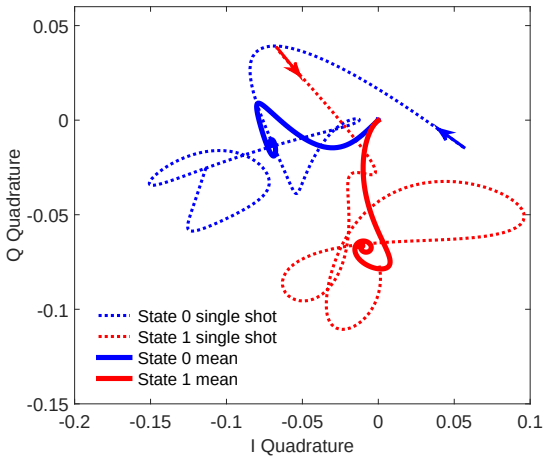

FIG. 1. Mean trajectories and single-shots for $|0\rangle$ (blue) and $|1\rangle$ (red) preparations (color online). The $|0\rangle$ ($|1\rangle$) single-shot trajectory (blue (red)-dotted) has arrow pointing up (down) and to the left (right). The mean trajectories of $|0\rangle$ (blue-solid) and $|1\rangle$) (red-solid) have steady-states of $\sim$ (-0.07,-0.02) and (-0.01,-0.07).

The current method of classifying trajectories [3] is to integrate each trajectory with a filter (kernel,weight) function $w(j)$. Formally, if $x \in \mathbb{R}^M$ is a single trajectory, under the assumption that the covariance matrices Σ_0 and Σ_1 of each class are equal, Gaussian, and diagonal, the optimal $w(j)$ is equal to the ratio of the difference between the mean trajectories μ_0 and μ_1 and the variance v. If Σ_d is the diagonal covariance matrix then

$$f(x) = x^T \left[\Sigma_d^{-1}(\mu_0 - \mu_1)\right]. \tag{2}$$

The assignment fidelity using this method is $\mathcal{F}_a = 0.9586$.

The "achievable" assignment fidelity for our experiment, $\mathcal{F}_{\text{ach}}$, is the fidelity we would obtain under ideal noise conditions. By "ideal noise" we mean the noise satisfies the assumptions above for the method of [3] to hold. The details of this calculation, along with a brief introduction to measurements and amplifiers, is contained in [16]. We obtain

$$\mathcal{F}_{\text{ach}} = 0.9999 \pm 0.0001, \tag{3}$$

and so there is a large discrepancy between $\mathcal{F}_{\text{ach}}$ and $\mathcal{F}_a$ that is due to a wide variety of factors such as state-preparation errors and non-Gaussian/non-linear effects. This discrepancy provides the motivation for investigating better methods for classifying trajectories.

The idea behind machine learning (ML) classification is to obtain a classifying (discriminating) surface in $\mathbb{R}^M$ under constraints such as the form of the noise. For Gaussian noise the optimal discriminator is a quadratic surface [18] (quadratic discriminant analysis-QDA) given by

$$f_{\text{QDA}}(x) = -\frac{1}{2}x^T\left[\Sigma_0^{-1} - \Sigma_1^{-1}\right]x + x^T\left[\Sigma_0^{-1}\mu_0 - \Sigma_1^{-1}\mu_1\right], \tag{4}$$

and the threshold is

$$v_{\text{QDA}} = \frac{1}{2}\left(\mu_0^T\Sigma_0^{-1}\mu_0 - \mu_1^T\Sigma_1^{-1}\mu_1\right) + \frac{1}{2}\log\left(\frac{|\Sigma_0|}{|\Sigma_1|}\right), \tag{5}$$

where $|\cdot|$ represents the determinant. If $\Sigma_0 = \Sigma_1$ the quadratic term in Eq. (4) disappears and the surface reduces to a hyperplane [4] (linear discriminant analysis-LDA),

$$f_{\text{LDA}}(x) = x^T\left[\Sigma^{-1}(\mu_0 - \mu_1)\right]. \tag{6}$$

Comparing with Eq. (2) we see the current method of [3] is equivalent to LDA with the added assumption of diagonal covariance matrices.

QDA can achieve a more accurate value of $\mathcal{F}_a$ as it allows for $\Sigma_0 \neq \Sigma_1$ and thus a quadratic (non-linear) discriminating surface. We computed $\mathcal{F}_a$ using the "fitcdiscr" function in Matlab for four different methods: LDAd, LDA, QDAd, and QDA ("d" represents diagonal covariance matrix and LDAd is the method of [3]). The results are in the second column of Table I. Not surprisingly, we find QDAd improves upon LDAd and allowing non-diagonal covariance matrices produces higher $\mathcal{F}_a$. The values in Table I are the sample means from 100 repetitions. The sample variances σ^2 are typically on the order of 1×10^{-8} indicating stable and reproducible results.

A value of $\mathcal{F}_a$ for QDA was not attainable due to singular covariance matrices, which is a result of overfitting the data (having more variables than required from the correlation time in the trajectories). To remedy this problem, we first perform dimensionality reduction using principal component analysis (PCA) [20] and find 99.9% of the variance in the data can be accounted for in a subspace of dimension ~ 20. Loosely speaking, this implies correlation times of $\sim 85 - 180$ ns. The results with a PCA pre-processing step (using "princomp" in Matlab) are in the third column of Table I. A value of QDA can now be computed and as expected it provides the highest $\mathcal{F}_a$ out of all cases considered.

TABLE I. Assignment fidelities for various discriminant analysis methods. See text for details.

Method	All time-points	PCA
LDAd	0.9586	0.9557
LDA	0.9701	0.9586
QDAd	0.9627	0.9648
QDA	–	0.9712

These classification methods have assumed Gaussian noise. More robust methods are needed as we expect non-Gaussian behavior. We approach this in two ways. The first is via the support vector machine (SVM) [5, 22], which makes no assumption on the form of the noise and can be extended to extremely general non-linear discriminating surfaces. The second is to utilize "clustering" methods in ML to naturally goup the data into clusters from which we perform *multi-class* classification. We first describe the SVM method.

The linear SVM is a quadratic program based on maximizing the minimum margins of the data, where the margin of a data point is its distance to a separating hyperplane [16]. The non-linear SVM is derived from the dual form of the linear SVM by defining a kernel that maps the data to a higher-dimensional space. The key is that the linear SVM in the higher-dimensional space allows for non-linear discrimination in $\mathbb{R}^M$. Due to its generality and simplicity, we chose a radial basis (Gaussian) function kernel. We implemented the SVM using the "fitcsvm" function in Matlab. The classification was repeated 100 times and the mean values with the optimal soft-margin parameter are contained in Table II (see [16] for details). The sample variances σ^2 in $\mathcal{F}_a$ are approximately 1.9×10^{-8} indicating stable results. The non-linear SVM produces the highest assignment fidelity out of all methods considered thus far, indicating non-linear effects are present.

Our second method for implementing a non-linear classifier combines classification and *clustering* algorithms. Clustering naturally groups the data into subsets and is "unsupervised" since it requires no training data. We utilize k-means clustering [6] since it has features that suit our purposes well, however we anticipate similar results can be obtained with other standard clustering algorithms such as heirarchical methods. For an explicit formulation of the k-means clustering problem see [16].

TABLE II. Assignment fidelities for SVM methods. See text for details.

Method	All time-points	PCA
Linear SVM	0.9753	0.9571
Non-linear SVM	0.9821	0.9739

We used the Matlab "kmeans" function to find $k = 3$ clusters in each of C_0 and C_1. We chose $k = 3$ to take into account both variance and systematic effects. The mean trajectories and size of the six subclasses are given in Fig. 2. We see C_0 is split relatively evenly into the subclasses $S_{0,1}$, $S_{0,2}$, $S_{0,3}$ that capture variance in the trajectories. We do not see a subclass of C_0 corresponding to heating of the ground state, however we implemented k-means for larger k and found a heating subclass of size ~ 230 for $k = 7$ (see Fig. 5).

C_1 has strikingly different properties as subclass $S_{1,2}$ is comprised of T_1 processes. $S_{1,1}$ and $S_{1,3}$ are similar in size and capture variance in the trajectories. The key point is we have found *explicit shot indices* for T_1 events. We verified that $S_{1,2}$ is comprised of T_1 trajectories by performing k-means with $k = 4$. We found that the $S_{1,1}$ and $S_{1,3}$ subclasses remain virtually the same while the T_1 subclass $S_{1,2}$ is now split into two according to variance in these trajectories (see Fig. 6). From Fig. 2, $\sim 9\%$ of the $|1\rangle$ preparations result in a T_1 event, which is consistent with the percentage calculated from system parameters [16], $1 - e^{-2.6/29} \sim 8.6\%$.

To perform classification, we lift the T_1 subclass $S_{1,2}$ to a class C_2 of its own, redefine $C_1 = S_{1,1} \cup S_{1,3}$, keep C_0 as before, and perform *multi-class* classification on C_0, C_1, and C_2. We implemented four multi-class algorithms in Matlab; multi-class LDA, multi-class SVM, "TotalBoost", and "RUSBoost". The latter two are examples of boosting algorithms which assemble an ensemble of weak learners (classifiers) in a network to create a final strong learner by iteratively re-weighting data points according to previous results [24]. The RUSBoost method [25] is particularly useful since it is tailored to the case of one class (here C_2) being significantly smaller than the rest.

The results are in Table III. We again see an increase in assignment fidelities over the discriminant analysis methods of Table I. Not surprisingly, RUSBoost provides the most significant increase. We repeated the k-means algorithm 50 times with random initializations and found it to be relatively stable (sample variance σ^2 of $\mathcal{F}_a \sim 3 \times 10^{-6}$). We repeated this using fixed initialization of the means and obtained a variance of 0.

Out of all methods considered, non-linear SVM's pro-

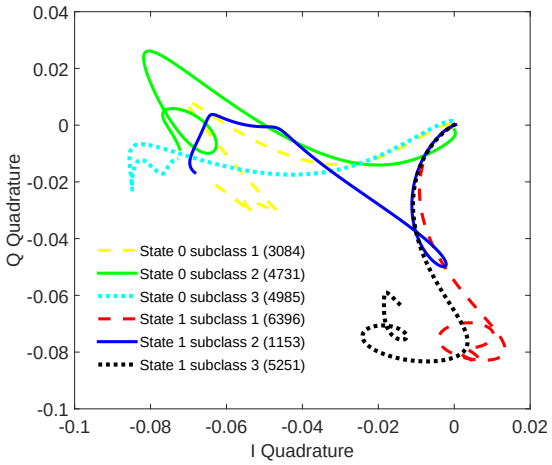

FIG. 2. Subclasses found from k-means algorithm (color online). C_0 and C_1 have three subclasses, the trajectory representing each subclass is the mean over all subclass trajectories. The subclasses of C_0 (yellow-dashed, green-solid, light blue-dotted) have paths that initially move up and left. The subclasses of C_1 (red-dashed, blue-solid, black-dotted) have paths that initially move down and right. The T_1 subclass (blue-solid) of C_1 initially moves down and right but abruptly changes its path to move up and left. The legend numbers are subclass sizes.

duce the greatest increase in $\mathcal{F}_a$ (0.9586 to 0.9821). We also note all methods are relatively stable with reproducible assignment fidelities (each method was repeated ~ 100 times; the sample means of $\mathcal{F}_a$ are the table values and the sample variances are $\sim 1 \times 10^{-8}$).

TABLE III. Assignment fidelities from multi-class classification. See text for details.

Method	All time-points	PCA
Multi-LDA	0.9768	0.9689
Multi-SVM	0.9784	0.9717
TotalBoost	0.9527	0.9413
RUSBoost	0.9788	0.9723

While we have improved $\mathcal{F}_a$ to 0.9821, we are still far from $\mathcal{F}_{\text{ach}} = 0.9999$. It is possible much of the remaining discrepancy comes from T_1 events. To investigate this we propose the simple error diagnosis test of replacing each T_1 event found from the k-means algorithm with a random element from $S_{1,2} \cup S_{1,3}$. This provides a measure of $\mathcal{F}_a$ when T_1 is negligible. The means of 100 samples for each method are contained in Table IV (variances are $\sim 1 \times 10^{-8}$). Non-linear SVM produces the highest value of $\mathcal{F}_a$ however *for all methods* $\mathcal{F}_a > 0.99$, which is more consistent with $\mathcal{F}_{\text{ach}} = 0.9999$. This confirms T_1 events are the significant reason for lower $\mathcal{F}_a$ values than expected.

One attempt to reduce the significance of T_1 is to reduce T, however this implies the trajectories will spend less time near their steady states and assignment errors due to variance will increase.To observe this, we truncated the trajectories to different T and calculated $\mathcal{F}_a$ using the non-linear SVM. From Fig. 3 we see $T = 2.6\mu$s appears close to optimal. Moreover, a much shorter measurement time of $\sim 1.2\mu$s (not shown in Fig. 3) is needed to achieve $\mathcal{F}_a$ from LDA. This is a strong message that better classifiers can allow for shorter measurement times. Longer measurement times than the current 2.6μs decrease $\mathcal{F}_a$ due to an increase in T_1 events.

TABLE IV. Assignment fidelities with replacement of T_1 events. See text for details..

Method	All time-points	PCA
LDAd	0.9920	0.9909
LDA	0.9921	0.9928
QDAd	0.9918	0.9908
QDA	–	0.9927
Linear SVM	0.9936	0.9943
Non-linear SVM	0.9945	0.9949

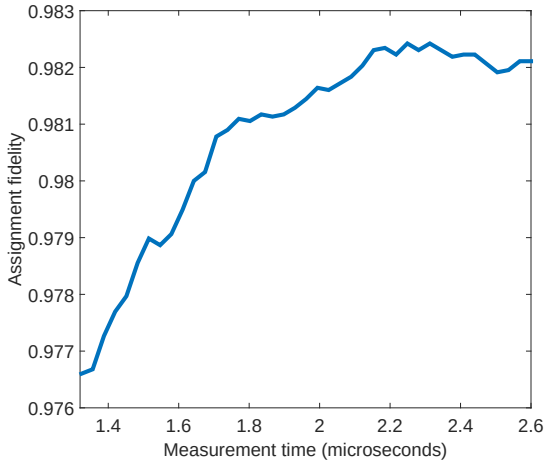

FIG. 3. Varying measurement time.

To conclude, we have utilized ML to understand and improve the readout in a superconducting system. We find more sophisticated classification algorithms can potentially allow for shorter measurement times and increase assignment fidelities. Non-linear SVM's provided the largest increase in assignment fidelity from 0.9586 to 0.9821 ($\sim 2.4\%$). Clustering helped diagnose the prevalence of systematic effects by finding clusters in the data corresponding to single-shot identification of heating and T_1 effects. We verified T_1 events are a significant source of error as the assignment fidelity increases from 0.9821 to 0.9945 when the T_1 cluster is replaced with typical trajectories. This is more consistent with our achievable fidelity and the remaining discrepancy can be due to ef-

fects such as heating and state-preparation errors. Moving forward, we expect these methods will help provide insight for improving readout, especially when non-linear and non-Gaussian effects are present.

We acknowledge support from ARO under contract W911NF-14-1-0124 and IARPA under contract W911NF-10-1-0324. We acknowledge helpful discussions with Oliver Dial, Stefan Filipp, Blake Johnson, Jim Rozen, Colm Ryan, Marcus Silva, and Matthias Steffen.

[1] P. Shor, in *Proceedings of the 37'th Annual Symposium on Foundations of Computer Science (FOCS)* (IEEE Press, Burlington, VT, 1996).
[2] H. Paik et al., Phys. Rev. Lett. **107**, 240501 (2011).
[3] J. B. Chang et al., Applied Physics Letters **103**, 012602 (2013).
[4] R. Barends et al., Physical Review Letters **111**, (2013).
[5] R. Barends et al., Nature **508**, 500 (2014).
[6] F. Mallet et al., Nat Phys **5**, 791 (2009).
[7] N. Bergeal et al., Nature **465**, 64 (2010).
[8] J. E. Johnson et al., Phys. Rev. Lett. **109**, 050506 (2012).
[9] D. Ristè et al., Phys. Rev. Lett. **109**, 050507 (2012).
[10] A. Blais et al., Phys. Rev. A **69**, 062320 (2004).
[11] J. Koch et al., Phys. Rev. A **76**, 042319 (2007).
[12] J. Gambetta et al., Phys. Rev. A **77**, 012112 (2008).
[13] A. D. Córcoles et al., arXiv:1410.6419 (2014).
[14] A. Hentschel and B. C. Sanders, Phys. Rev. Lett. **104**, 063603 (2010).
[15] M. Gu and W. Kotowski, New Journal of Physics **12**, 123032 (2010).
[16] See Supplemental Material for further details.
[17] C. A. Ryan et al., arXiv:1310.6448 (2013).
[18] T. M. Cover, Electronic Computers, IEEE Transactions on **EC-14**, 326 (1965).
[19] R. A. Fisher, Annals of Eugenics **7**, 179 (1936).
[20] K. Pearson, Philosophical Magazine Series 6 **2**, 559 (1901).
[21] C. Cortes and V. Vapnik, Machine Learning **20**, 273 (1995).
[22] B. E. Boser, I. M. Guyon, and V. N. Vapnik, in *Proceedings of the Fifth Annual Workshop on Computational Learning Theory* (ACM, New York, NY, USA, 1992), COLT '92, pp. 144–152.
[23] S. Lloyd, Information Theory, IEEE Transactions on **28**, 129 (1982).
[24] C. Bishop, *Pattern Recognition and Machine Learning* (Springer, 2007).
[25] C. Seiffert et al., Systems, Man and Cybernetics, Part A: Systems and Humans, IEEE Transactions on **40**, 185 (2010).

SUPPLEMENTAL MATERIAL

Measurement and linear amplification in circuit QED

In the dispersive regime of circuit-QED the resonator frequency depends on the qubit state so that driving the cavity and observing the output of the cavity corresponds to a quantum non-demolition measurement. For outcome "0" (1) the mean voltage leaving the cavity represents a time-dependent coherent state, denoted $\alpha_{0(1)}(t)$, that is typically small in magnitude. The phase-space evolution of $\alpha_{0(1)}(t)$ is determined by the deterministic differential equation

$$\dot{\alpha}_{0(1)} = -i\mathcal{E}(t) - i\left(\omega_r + \chi_{0(1)}\right)\alpha_{0(1)} - \left(\kappa\alpha_{0(1)}\right)/2, \quad (7)$$

where $\mathcal{E}(t) = \mathcal{E}_x(t)\cos(\omega_m t) + \mathcal{E}_y(t)\sin(\omega_m t)$ is the measurement drive and κ is the cavity decay rate.

For our experiment, the qubit transition frequency is $\omega/2\pi = 5.415$ GHz and the readout resonator frequency is $\omega_R/2\pi = 6.693$ GHz. The qubit anharmonicity is $\sim$ 330 MHz with T_1 (energy relaxation) time of 29μs and coherence time $T_2^{\text{echo}} = 22\mu$s. The dispersive shift and line width of the readout resonator are measured to be $2\chi/2\pi = -2.8$ MHz and $\kappa/2\pi = 1210$ kHz, respectively so that $\chi \sim \kappa$.

The output mode $b(t)$ from the resonator can be related to the field inside the resonator by

$$\langle b(t)\rangle = \sqrt{\kappa}\frac{\beta(t)}{2}\langle z\rangle + \sqrt{\kappa}\frac{\nu(t)}{2}, \quad (8)$$

where $\beta(t) = \alpha_0(t) - \alpha_1(t)$ is the separation between the pointer states, $\nu(t) = \alpha_0(t) + \alpha_1(t)$ is the mean value of the coherent states and $\langle z\rangle = p_0 - p_1$ represents the qubit information with p_0 and p_1 being the probability of the qubit being in state 0 and 1 respectively [1]. Typically this field is amplified by a linear phase-preserving amplifier to the output mode $c(t)$,

$$c(t) = \sqrt{G}b(t) + h(t), \quad (9)$$

where G is the power gain and $h(t)$ is the extra noise added by the amplifier. Here we have made an assumption that the bandwidth of the amplifier is constant and larger than the bandwidth of the signal being measured. Since the commutators must satisfy

$$\left[c(t), c^\dagger(t')\right] = \left[b(t), b^\dagger(t')\right] = \mathbb{1}\delta(t-t'), \quad (10)$$

the noise must satisfy

$$\left[h(t), h^\dagger(t')\right] = \mathbb{1}(1-G)\delta(t-t'), \quad (11)$$

which, from the generalized uncertainty principle of an arbitrary operator R,

$$|\Delta R|^2 \geq \frac{1}{2}|\langle[R, R^\dagger]\rangle| \quad (12)$$

implies $|\Delta h|^2 \geq \frac{1}{2}|1-G|\delta(t-t')$ and $|\Delta b|^2 \geq \frac{1}{2}\delta(t-t')$. Using Eq. 9 the noise in the output mode $c(t)$ is

$$|\Delta c|^2 \geq \frac{G(1+2A)\delta(t-t')}{2}, \tag{13}$$

where

$$A \geq \frac{1}{2}\left(1-\frac{1}{G}\right) \tag{14}$$

is the added noise normalized by the gain [2]. In our view this is the best number to quantify an amplifier and for the quantum limit it takes the value 1/2 for a phase preserving amplifier. Other useful quantities are the instantaneous input and output signal-to-noise ratios of the amplifier defined by

$$\mathrm{SNR}_{\mathrm{out}} = \frac{|\langle c\rangle|^2}{|\Delta c|^2} = \eta_l |\langle b\rangle|^2 \delta t \tag{15}$$

$$\mathrm{SNR}_{\mathrm{in}} = \frac{|\langle b\rangle|^2}{|\Delta b|^2} = 2|\langle b\rangle|^2 \delta t, \tag{16}$$

where $\eta_l = 2/(1+2A) \leq 1$ is the efficiency of the amplifier and represents how well the input field is mapped to the output field. Another useful quantity is the noise-figure f of the amplifier, which is the ratio $\mathrm{SNR}_{\mathrm{in}}/\mathrm{SNR}_{\mathrm{out}}$. We see that $f = 1+2A$ which is 2 for a quantum limited amplifier.

In circuit QED the information about the qubit state is contained in a single quadrature. From Eq. 8 this is the quadrature set by $\theta = \arg(\beta(t))$ and as a result when we subtract the mean value we obtain

$$\mathrm{SNR}_{\mathrm{out}} = \frac{\kappa\eta_l|\beta(t)|^2\delta t}{4}. \tag{17}$$

This overestimates the noise as the information is only in one quadrature. Defining the measurement quadrature $I(t) = \mathrm{Re}[c(t)e^{-i\theta}]$ it is simple to show that

$$|\Delta I|^2 \geq \frac{1}{2}|\Delta c|^2 = \frac{G}{4}(1+2A)\delta(t-t'), \tag{18}$$

giving an instantaneous SNR_I

$$\mathrm{SNR}_I = \frac{\langle I(t)\rangle^2}{|\Delta I|^2} = \eta\kappa|\beta(t)|^2\delta t, \tag{19}$$

where $\eta = 1/(1+2A) \leq \frac{1}{2}$ is the efficiency of measuring information in a single quadrature for a linear phase preserving amplifier. Note this is a factor of two less than the efficiency of the amplifier.

A general linear amplifier can be described by the output mode $c(t)$,

$$c(t) = Mb(t) + Lb^\dagger(t) + h(t), \tag{20}$$

and from preservation of the commutation relations

$$\left[h(t), h^\dagger(t')\right] = \mathbb{1}(1-M^2+|L|^2)\delta(t-t'), \tag{21}$$

we have $|\Delta h|^2 \geq \frac{1}{2}|(1-M^2+|L|^2)|\delta(t-t')$. Setting $|L| = \sqrt{M^2-1}$ results in $|\Delta h|^2 \geq 0$. To amplify a single quadrature $I = \mathrm{Re}[c(t)e^{-i\theta}]$ (phase sensitive amplifier) we set $L = \sqrt{M^2-1}e^{2i\theta}$ giving

$$I(t) = \sqrt{G}\mathrm{Re}[b(t)e^{-i\theta}] + \mathrm{Re}[h(t)e^{-i\theta}], \tag{22}$$

$$Q(t) = \frac{1}{\sqrt{G}}\mathrm{Im}[b(t)e^{-i\theta}] + \mathrm{Im}[h(t)e^{-i\theta}], \tag{23}$$

where $G = (M+\sqrt{M^2-1})^2$ is the power gain. From the generalized uncertainty principle

$$|\Delta I|^2 \geq \frac{G_s}{4}(1+2A_s)\delta(t-t'), \tag{24}$$

where $A_s \geq 0$ is the gain normalized added noise. The instantaneous SNR_I for the quadrature I is

$$\mathrm{SNR}_I = \frac{\langle I(t)\rangle^2}{|\Delta I|^2} = \eta\kappa|\beta(t)|^2\delta t. \tag{25}$$

where $\eta = 1/(1+2A_s) \leq 1$. That is, by using a phase sensitive amplifier tuned to the correct phase the effective SNR can be a factor of two better then a phase preserving amplifier.

Filtering protocol for ideal noise and the achievable fidelity

In a typical measurement protocol the measurement outcome is the integration of the signal $c(t)$ from the amplifier with a weighting kernel (filter),

$$S = \int_0^{t_m} \mathrm{Re}[w(t)c(t)]dt = \sum_j |w_j|\mathrm{Re}[e^{-i\phi_j}c_j]\delta t, \tag{26}$$

where t_m is the measurement time and $w_j = |w_j|e^{-i\phi_j}$ is the kernel. Under the assumption that the noise is symmetric a useful measure to quantify the measurement is the separation

$$R = \frac{(\langle S_0\rangle - \langle S_1\rangle)^2}{\mathrm{var}(S)}, \tag{27}$$

where 0 and 1 label the two states of the qubit. In Ref. [3] it was shown that the optimal kernel under the additional assumptions of Gaussian and diagonal covariance matrices is found by maximizing R and is given by

$$w_j = \frac{\beta_j^*}{\mathrm{var}(I_j)}, \tag{28}$$

where β_j is the separation of the pointer states and $I_j = \mathrm{Re}[e^{-i\theta_j}c_j]$.

The achievable fidelity $\mathcal{F}_{\mathrm{ach}}$ is the value of the assignment fidelity $\mathcal{F}_a$ assuming that the noise in our system is

ideal (symmetric and diagonal Gaussian covariance matrices). In reality, the noise does not satisfy these properties so we must fit it to an ideal noise model. To do this, we implemented the above filtering method on our data and plotted the resulting distributions for the different classes ($|0\rangle$ and $|1\rangle$) of S in Fig. 4. If the noise was ideal we should obtain two Gaussian histograms with identical variance. In reality, we see there are significant non-Gaussian statistics and so we fit the histograms to double Gaussian distributions of equal variance to obtain the means and standard deviation we would expect in the ideal noise case. We obtain $\langle S_0\rangle = -0.2149 \pm 0.0149$ and $\langle S_1\rangle = 2.750 \pm 0.0161$, and a standard deviation of $\sqrt{\text{var}(S)} = 0.6758 \pm 0.0109$, which gives a value for R (defined in Eq. 27) of 55.56 ± 1.02.

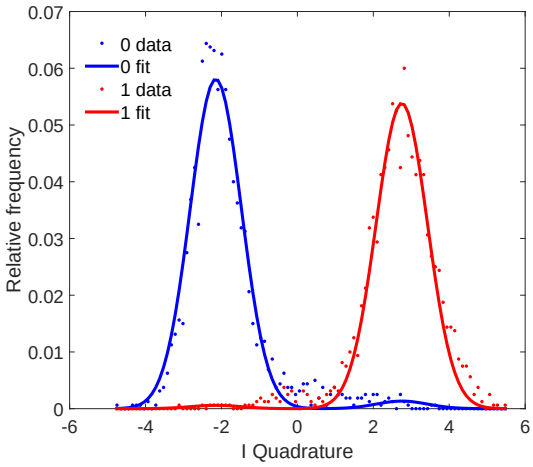

FIG. 4. Projected data and double Gaussian fits of the $|0\rangle$ (blue-left) and $|1\rangle$ (red-right) preparation classes (color online).

We can now compute $\mathcal{F}_{\text{ach}}$ by combining the ideal noise assumption with the definition of the assignment fidelity,

$$\mathcal{F}_a = 1 - \left(\mathbb{P}\left[0|1\right] + \mathbb{P}\left[1|0\right]\right)/2. \tag{29}$$

Here $\mathbb{P}[0|1]$ ($\mathbb{P}[1|0]$) is the probability of obtaining outcome "1" ("0") given the system was prepared in $|0\rangle$ ($|1\rangle$). Since the noise is ideal $\mathcal{F}_{\text{ach}}$ takes the form

$$\mathcal{F}_a = 1/2 + \text{erf}(\sqrt{R/8})/2, \tag{30}$$

where R is defined in Eq. 27. From this expression we obtain

$$\mathcal{F}_{\text{ach}} = 0.9999 \pm 0.0001 \tag{31}$$

for our system. We note that the separation R in the time independent limit can be related to the signal-to-noise defined above by $\text{SNR}_I = R/4$. This is straightforward to show by direct substitution of the quantities in Eq. 27. R is used here as it is a standard measure of separation in ML [4].

SVM's and k-means clustering

The quadratic program for the SVM [5] is given by the equation

$$\begin{aligned} \underset{w,b}{\text{minimize}} \quad & \|w\|^2/2 \\ \text{subject to} \quad & y^{(i)}(w^T x^{(i)} + b) \geq 1, \;\; i \in \{1, ..., M\}, \end{aligned} \tag{32}$$

where y is the expected outcome (taken to be -1 or 1). This has a quadratic objective function with linear inequality constraints. The soft-margin formalism adds slack variables representing the degree of misclassification to the constraints in Eq. (32) and modifies the objective function to include a mislabelling cost term. The modified quadratic program that includes a soft-margin is given by the equation

$$\begin{aligned} \underset{w,b,\xi}{\text{minimize}} \quad & \|w\|^2/2 + C\sum_{j=1}^{M} \xi_j \\ \text{subject to} \quad & y^{(i)}(w^T x^{(i)} + b) \geq 1 - \xi_i, \;\; i \in \{1, ..., M\}. \end{aligned} \tag{33}$$

In the dual form of this problem, the ξ variables vanish and C is a "box constraint" that bounds the Lagrange multipliers. We implemented 10-fold cross-validation on the training set and found a misclassification error of 0.0163 which agrees well with $\mathcal{F}_a = 0.9825$. This implies large C value is likely not needed and varying C between 0 and 100 gave an optimal value of $\mathcal{F}_a$ close to 1.

k-means clustering [6] is formulated as the following optimization problem

$$\underset{\mathbf{S}}{\text{minimize}} \quad \sum_{j=1}^{k} \sum_{x \in S_j} \|x - \mu_j\|^2. \tag{34}$$

Here there are k sets $S = \{S_1, ..., S_k\}$ with means μ_j and so the goal is to partition the data into k sets that minimize the within-set distance to the mean μ_j. The k-means algorithm only has guaranteed convergence to a local minimum and finding the global optimum is an NP-hard problem [7]. Therefore initialization of the subclass means in Eq. 34 can be important to find meaningful solutions. Typically initialization is a random process that can lead to small variation in the solutions. This can be circumvented by explicitly defining the initial means to be the average of the output subclass means over many realizations. This helps ensure the algorithm is reproducible and more stable.

The plot of the heating subclass of C_0 found from the k-means algorithm with $k = 7$ is given in Fig. 5 and the plot of the four subclasses of C_1 found from the k-means algorithm with $k = 4$ is given in Fig. 6.

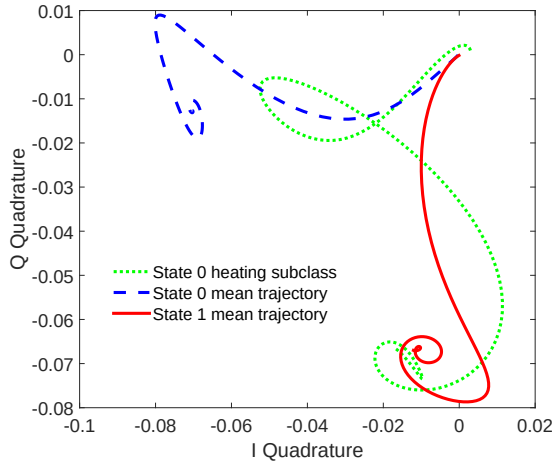

FIG. 5. Heating subclass of C_0 (green-dotted) found from a k-means algorithm with $k = 7$ superimposed on means of $|0\rangle$ (blue-dashed) and $|1\rangle$ (red-solid) classes (color online).

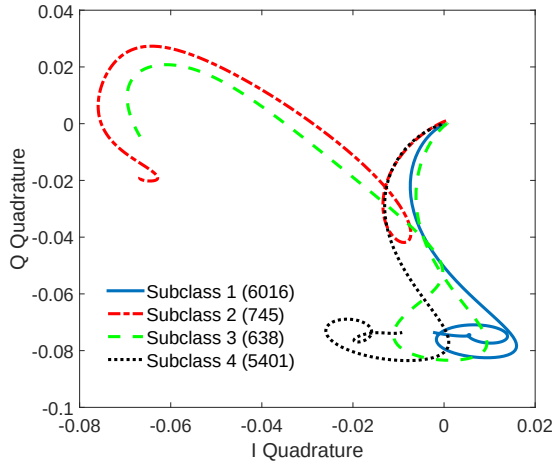

FIG. 6. Four subclasses of C_1 found from a k-means algorithm. There are two T_1 subclasses (green-dashed and red-dash-dotted) that split the subclass found from $k = 3$. The other subclasses (black-dotted and blue-solid) are comprised of typical trajectories of a $|1\rangle$ state (color online).

[1] J. Gambetta et al., Phys. Rev. A **77**, 012112 (2008).
[2] C. M. Caves, Phys. Rev. D **26**, 1817 (1982).
[3] C. A. Ryan et al., arXiv:1310.6448 (2013).
[4] R. A. Fisher, Annals of Eugenics **7**, 179 (1936).
[5] C. Cortes and V. Vapnik, Machine Learning **20**, 273 (1995).
[6] S. Lloyd, Information Theory, IEEE Transactions on **28**, 129 (1982).
[7] D. Aloise et al., Machine Learning **75**, 245 (2009).